Docket No: 40291/2000100

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In the Application of:

Ryszard BLESZYNSKI and

Man D. TRINH

Serial No.: To Be Assigned

Filing Date: October 23, 2000

For: MEMORY MANAGEMENT SYSTEM AND ALGORITHM FOR NETWORK PROCESSOR ARCHITECTURE

Examiner: To Be Assigned

Group Art Unit: To Be Assigned

PRELIMINARY AMENDMENT

Attention: Office of Initial Patent Examination
Assistant Commissioner for Patents
Washington, D.C. 20231

Dear Sir:

Prior to examination on the merits, please enter the following amendments.

In the Specification:

On page 2 (page 2 begins the "Background of the Invention" section), after line 2 and before the first paragraph, please add the following subsection heading (i.e., the "Field of the Invention" is a subsection of the "Background of the Invention" section):

"**Field of the Invention**"

On page 2, after line 8 and before the second paragraph, please add the following subsection heading:

"**Background Information**"

On page 2, please delete line 20 (i.e., please delete the words "Figure 1").

On page 2, line 21, please delete Figure 1 entitled "Network Processor in a Line Card".

On page 3, please delete line 20 (i.e., please delete the words "Figure 2").

On page 3, line 21, please delete Figure 2 entitled "Moore's Law Vs Bandwidth Demand".

On page 5, please delete line 1 (i.e., please delete the words "Figure 3").

On page 5, line 2, please delete Figure 3 entitled "Memory Hierarchy in a Computer System".

On page 5, after the second paragraph that begins with "Caching theory works well...." and before the "Summary of the Invention" section, please add the following new paragraph:

Therefore, it is desirable to have a system and method to efficiently access a memory unit while processing network traffic.

On page 6, please delete line 12 (i.e., please delete the words "Figure 4").

On page 6, line 13, please delete Figure 4 entitled "Memory Management Subsystems".

On page 8, after line 5, please add the following section and the brief descriptions of the figures:

Brief Description of the Drawings

Figure 1 illustrates a prior art line card and its components.

Figure 2 illustrates Moore's Law versus the Internet bandwidth demand curve.

Figure 3 illustrates a prior art multilevel memory hierarchy within a processor.

Figure 4 illustrates an embodiment of a memory management subsystem according to the present invention.

Figure 5 illustrates an embodiment of a network processor according to the present invention.

Figure 6 illustrates an embodiment of a payload channel sequence table according to the present invention.

Figure 7 illustrates an embodiment of a bandwidth balancing flowchart according to the present invention.

On page 10, please delete line 7 (i.e., please delete the words "Figure 5").

On page 10, line 8, please delete Figure 5 entitled "Bandwidth Balancer Data Structure".

On page 12, please delete line 5 (i.e., please delete the words "Figure 6").

On page 12, line 6, please delete Figure 6 entitled "Payload Channel Sequence Table Structure".

On page 14, please delete line 1 (i.e., please delete the words "Figure 7").

On page 14, line 2, please delete Figure 7 entitled "Bandwidth Balancing Algorithm Flow Chart".

After the claims, please add the following Abstract section:

Abstract of the Disclosure

An embodiment of this invention pertains to a system and method for balancing memory accesses to a low cost memory unit in order to sustain and guarantee a desired line rate regardless of the incoming traffic pattern. The memory unit may include, for example, a group of dynamic random access memory units. The memory unit is divided into memory channels and each of the memory channels is further divided into memory lines, each of the memory lines includes one or more buffers that correspond to the memory channels. The determination as to which of one or more buffers within a memory line an incoming information element is stored is based on factors such as the number of buffers pending to be read within each of the memory channels, the number of buffers pending to be written within each of the memory channels, and the number of buffers within each of the memory channels that has data written to it and is waiting to be read.

In the Claims:

Please add the following claims:

1. (new) A method to optimally access a memory unit where the memory unit is logically partitioned to form a plurality of memory channels, the plurality of memory channels are further logically partitioned to form a plurality of memory lines, each of the plurality of memory lines includes a plurality of buffers and each of the plurality of buffers corresponds to a separate one of the plurality of memory channels, comprising:

determining at least one load value of each of the plurality of memory channels; and

based on the determined at least one load value, selecting a particular one of the plurality of memory channels.

2. (new) The method of claim 1 wherein the step of determining the at least one load value of each of the plurality of memory channels includes determining, for each of the plurality of memory channels, the number of pending read requests.

3. (new) The method of claim 1 wherein the step of selecting the particular one of the plurality of memory channels includes selecting the particular one of the plurality of memory channels that has a lowest number of pending read requests.

4. (new) The method of claim 1 wherein the step of determining the at least one load value of each of the plurality of memory channels includes determining, for each of the plurality of memory channels, at least one of the number of pending write requests, and the number of active buffers which is the number of a particular one of the plurality of buffers that is unavailable and corresponds to the particular one of the plurality of memory channels in each of the plurality of memory lines.

5. (new) The method of claim 1 wherein the step of selecting the particular one of the plurality of memory channels includes selecting the particular one of the plurality of memory channels that has at least one of a lowest number of pending write requests, a lowest number of active buffers, and a corresponding channel identification number that is next in a round robin scheme.

6. (new) The method of claim 1 wherein the memory unit is a plurality of dynamic random access memory units.

7. (new) The method of claim 1 wherein each of the plurality of buffers has a fixed-size.

8. (new) The method of claim 7 further comprising

receiving an incoming information element;

if the size of the information element is greater than the fixed-size of each of the plurality of buffers, dividing the information element into a plurality of information element segments, each of the plurality of information element segments having a size less than or equal to the fixed-size of each of the at least one buffer; and

storing at least one of the information element and a particular one of the plurality of information element segments within a particular one of the plurality of buffers corresponding to the selected one of the plurality of memory channels at a particular one of the plurality of memory lines.

9. (new) The method of claim 1 wherein each of the plurality of memory channels has a width equal to a width of the memory unit divided by the number of the plurality of memory channels.

10. (new) A method to optimally access a memory unit where the memory unit is logically partitioned to form a plurality of memory channels, the plurality of memory channels are further logically partitioned to form a plurality of memory lines, each of the plurality of memory lines includes a plurality of buffers and each of the plurality of buffers corresponds to a separate one of the plurality of memory channels, comprising:

determining at least one load value of each of the plurality of memory channels; and

selecting a particular one of the plurality of memory channels that has a particular one of the at least one load value that is the lowest.

11. (new) The method of claim 10 wherein the step of determining the at least one load value of each of the plurality of memory channels includes determining, for each of the plurality of memory channels, at least one of the number of pending read requests, the number of pending write requests, and the number of active buffers which is the number of a particular one of the plurality of buffers that is unavailable and corresponds to the particular one of the plurality of memory channels in each of the plurality of memory lines.

12. (new) The method of claim 10 wherein the step of selecting the particular one of the plurality of memory channels that has the lowest determined load includes selecting the particular one of the plurality of memory channels that has at least one of a lowest number of pending read requests, a lowest number of pending write requests, a lowest number of active buffers, and a corresponding channel identification number that is next in a round robin scheme.

13. (new) The method of claim 10 wherein the memory unit is a plurality of dynamic random access memory units.

14. (new) The method of claim 10 wherein each of the plurality of buffers has a fixed-size.

15. (new) The method of claim 14 further comprising

receiving an incoming information element;

if the size of the information element is greater than the fixed-size of each of the plurality of buffers, dividing the information element into a plurality of information element segments, each of the plurality of information element segments having a size less than or equal to the fixed-size of each of the at least one buffer; and

storing at least one of the information element and a particular one of the plurality of information element segments within a particular one of the plurality of buffers corresponding to the particular one of the plurality of buffers to the selected one of the plurality of memory channels at a particular one of the plurality of memory lines.

16. (new) The method of claim 14 wherein each of the plurality of memory channels has a width equal to a width of the memory unit divided by the number of the plurality of memory channels.

17. (new) A method to optimally access a single hierarchical level memory unit, where the memory unit is logically partitioned to form a plurality of memory channels, the plurality of memory channels are further logically partitioned to form a plurality of memory lines, each of the plurality of memory lines includes a plurality of buffers and each of the plurality of buffers corresponds to a separate one of the plurality of memory channels, comprising:

determining, for each of the plurality of memory channels, at least one of the number of pending read requests, the number of pending write requests, and the number of active buffers which is the number of a particular one of the plurality of buffers that is unavailable and corresponds to the particular one of the plurality of memory channels in each of the plurality of memory lines; and

selecting a particular one of the plurality of memory channels that has at least one of a lowest number of pending read requests, a lowest number of pending write requests, a lowest

number of active buffers, and a corresponding channel identification number that is next in a round robin scheme.

18. (new) The method of claim 17 wherein each of the plurality of buffers has a fixed-size.

19. (new) The method of claim 18 further comprising

receiving an incoming information element;

if the size of the information element is greater than the fixed-size of each of the plurality of buffers, dividing the information element into a plurality of information element segments, each of the plurality of information element segments having a size less than or equal to the fixed-size of each of the at least one buffer; and

storing at least one of the information element and a particular one of the plurality of information element segments within a particular one of the plurality of buffers corresponding to the selected one of the plurality of memory channels at a particular one of the plurality of memory lines.

20. (new) The method of claim 17 wherein the single hierarchical level memory unit is a plurality of dynamic random access memory units.

21. (new) The method of claim 19 wherein the step of selecting the particular one of the plurality of memory channels includes

finding a first subset of the plurality of memory channels that is available at the particular one of the plurality of memory lines and has a lowest number of the pending read requests;

if the number of memory channels within the first subset of the plurality of memory channels equals one, setting the selected one of the plurality of memory channels to the first subset of the plurality of memory channels;

if the number of memory channels within the first subset of the plurality of memory channels is greater than one, then finding a second subset of the plurality of memory channels within the first subset of the plurality of memory channels that has the lowest number of the pending write requests;

if the number of memory channels within the second subset of the plurality of memory channels equals one, setting the selected one of the plurality of memory channels to the second subset of the plurality of memory channels;

if the number of memory channels within the second subset of the plurality of memory channels is greater than one, then finding a third subset of the plurality of memory channels within the second subset of the plurality of memory channels that has the lowest number of active buffers;

if the number of memory channels within the third subset of the plurality of memory channels equals one, setting the selected one of the plurality of memory channels to the third subset of the plurality of memory channels; and

if the number of memory channels within the third subset of the plurality of memory channels is greater than one, setting the selected one of the plurality of memory channels to a particular one of the third subset of the plurality of memory channels that has a corresponding channel identification number that is next in a round robin scheme.

22. (new) The method of claim 19 wherein the step of selecting the particular one of the plurality of memory channels includes

finding a first subset of the plurality of memory channels that has a lowest number of the pending read requests;

if the number of memory channels within the first subset of the plurality of memory channels equals one,

determining if the first subset of the plurality of memory channels at the particular one of the plurality of memory lines is available;

if the first subset of the plurality of memory channels at the particular one of the plurality of memory lines is available, setting the selected one of the plurality of memory channels to the first subset of the plurality of memory channels; and

if the first subset of the plurality of memory channels at the particular one of the plurality of memory lines is not available,

determining if at least one of the information element and the particular one of the plurality of information element segments can be stored within any

remaining one of the plurality of memory channels at the particular one of the plurality of memory lines without overloading that memory channel;

if at least one of the information element and the particular one of the plurality of information element segments can be stored within any remaining one of the plurality of memory channels, finding a second subset of the plurality of memory channels that has a next lowest number of the pending read requests; and

if at least one of the information element and the particular one of the plurality of information element segments cannot be stored within any remaining one of the plurality of memory channels, fetching a new one of the plurality of memory lines; and

if the number of memory channels within the first subset of the plurality of memory channels is greater than one, setting the selected one of the plurality of memory channels to a particular one of the first subset of the plurality of memory channels that has at least one of a lowest number of pending write requests, a lowest number of active buffers, and a corresponding channel identification number that is next in a round robin scheme.

23. (new) The method of claim 19 further comprising, upon storing at least one of the information element and the particular one of the plurality of information element segments within the particular one of the plurality of buffers corresponding to the selected one of the plurality of memory channels, setting a particular one of a plurality of payload channel occupancy bits that corresponds to the selected one of the plurality of memory channels.

24. (new) The method of claim 19 further comprising, reading the plurality of payload channel occupancy bits to determine if a corresponding one of the plurality of memory channels is available.

25. (new) The method of claim 19 further comprising, upon storing at least one of the information element and the particular one of the plurality of information element segments within the particular one of the plurality of buffers corresponding to the selected one of the plurality of memory channels at the particular one of the plurality of memory lines, writing a channel identification number corresponding to the selected one of the plurality of memory

channels to a buffer location field within a payload channel sequence table that corresponds to the particular one of the plurality of buffers.

26. (new) The method of claim 25 further comprising, upon storing at least one of the information element and the particular one of the plurality of information element segments within the particular one of the plurality of buffers corresponding to the selected one of the plurality of memory channels at the particular one of the plurality of memory lines, setting a value field within the payload channel sequence table that corresponds to the particular one of the plurality of buffers.

27. (new) The method of claim 26 further comprising, upon storing at least one of the information element and the particular one of the plurality of information element segments within the particular one of the plurality of buffers corresponding to the selected one of the plurality of memory channels at the particular one of the plurality of memory lines, if the data within the particular one of the plurality of buffers signals an end-of-packet, setting an end-of-packet field corresponding to the particular one of the plurality of buffers within the payload channel sequence table.

28. (new) The method of claim 27 further comprising fetching at least one of the information element and a portion of the information element by

determining at least one memory channel that stores at least one of the information element and the portion of the information element by reading the buffer location field corresponding to each of the plurality of buffers at a particular one of the plurality of memory lines until an end-of-packet field corresponding to that buffer signals the end-of-packet; and

reading the contents of each of an at least one buffer of the plurality of buffers at a particular one of the plurality of memory lines corresponding to each of the at least one memory channel.

29. (new) A system to optimally access a memory unit, comprising:

the memory unit that is logically partitioned to form a plurality of memory channels;

a traffic analyzer to determine at least one load of each of the plurality of memory channels; and

a bandwidth balancer to select a particular one of the plurality of memory channels based on the determined at least one load.

30. (new) The system of claim 29 wherein the plurality of memory channels of the memory unit are further logically partitioned to form a plurality of memory lines, each of the plurality of memory lines includes a plurality of buffers and each of the plurality of buffers corresponds to a separate one of the plurality of memory channels.

31. (new) The system of claim 29 further comprising

a plurality of write payload channel queues, each of the plurality of write payload channel queues corresponds to a separate one of the plurality of memory channels, each of the plurality of write payload channel queues stores at least one of an information element and a particular one of the information element segments to be written to a corresponding one of the plurality of memory channels;

a plurality of write request queues, each of the plurality of write request queues corresponds to a separate one of the plurality of write payload channel queues, a particular one of the plurality of write request queues stores a request to write the data within a corresponding one of the plurality of write payload channel queues to a corresponding one of the plurality of memory channels of the memory unit;

a plurality of read payload channel queues, each of the plurality of read payload channel queues corresponds to a separate one of the plurality of memory channels, each of the plurality of read payload channel queues stores at least one of an information element and a particular one of the information element segments that is retrieved from the memory unit; and

a plurality of read request queues, each of the plurality of read request queues corresponds to a separate one of the plurality of read payload channel queues, a particular one of the plurality of read request queues stores a request to retrieve from a corresponding one of the plurality of memory channels of the memory unit at least one of the information element and the particular one of the information element segments and store it in a corresponding one of the plurality of read payload channel queues.

32. (new) The system of claim 31 wherein the traffic analyzer includes

a plurality of pending write request counters to measure write request loads on the plurality of channels, each of the plurality of pending write request counters corresponds to a separate one of the plurality of write request queues;

a plurality of pending read request counters to measure read request loads on the plurality of channels, each of the plurality of pending read request counters corresponds to a separate one of the plurality of read request queues; and

a plurality of active buffer counters to measure stored data loads on the plurality of channels, each of the plurality of active buffer counters corresponds to a separate one of the plurality of write request queues that in turn corresponds to a particular one of the plurality of memory channels and each of the plurality of active buffer counters also corresponds to a separate one of the plurality of write request queues that in turn corresponds to the particular one of the plurality of memory channels.

33. (new) The system of claim 32 wherein

a particular one of the plurality of pending write request counters is incremented upon a corresponding one of the plurality of write request queues receiving a write request and decremented upon extracting the write request from the corresponding one of the plurality of write request queues;

a particular one of the plurality of pending read request counters is incremented upon a corresponding one of the plurality of read request queues receiving a read request and decremented upon extracting the read request from the corresponding one of the plurality of read request queues; and

a particular one of the plurality of active buffer counters is incremented upon a corresponding one of the plurality of write request queues receiving the write request and decremented upon a corresponding one of the plurality of read request queues receiving the read request.

34. (new) The system of claim 29 further comprising a payload channel occupancy vector, each element of the payload channel occupancy vector corresponds to a separate one of the plurality

of buffers at a particular one of the plurality of memory lines and each element of the payload channel occupancy vector indicates if a corresponding one of the plurality of memory buffers is available.

35. (new) The system of claim 29 further comprising a payload channel sequence table to specify an at least one memory channel of the plurality of memory channels at which at least one of the information element and a portion of the information element is stored.

36. (new) The system of claim 35 wherein the payload channel sequence table is partitioned to form a plurality of columns, each of the plurality of columns corresponds to a separate one of the plurality of memory channels, the plurality of columns are further partitioned to form a plurality of rows, each of the plurality of rows includes a plurality of buffer information units and each of the plurality of buffer information units includes

a buffer location field that specifies a particular one of the plurality of memory channels at which a particular one of the plurality of buffers at a particular one of the plurality of memory lines stores at least one of the information element and a particular one of the plurality of information element segments;

a value field that indicates whether the particular one of the plurality of buffers corresponding to the particular one of the plurality of memory channels at the particular one of the plurality of memory lines stores any data within that buffer; and

an end-of-packet field that indicates whether the particular one of the plurality of buffers corresponding to the particular one of the plurality of memory channels at the particular one of the plurality of memory lines stores data that signals an end-of-packet.

37. (new) The system of claim 29 wherein the memory unit is a plurality of dynamic random access memory units.

38. (new) The system of claim 30 further comprising a buffer management unit to provide a pointer to a new one of the plurality of memory lines.

39. (new) The system of claim 29 wherein each of the plurality of buffers has a length that is a fixed-size.

40. (new) The system of claim 39 wherein each of the plurality of memory channels has a width that is the fixed-size.

41. (new) A system to optimally access a memory unit, comprising:

the memory unit that is logically partitioned to form a plurality of memory channels;

a bandwidth management unit that includes

a traffic analyzer to determine at least one load of each of the plurality of memory channels; and

a bandwidth balancer to select a particular one of the plurality of memory channels based on the determined at least one load; and

a policy control unit to provide at least one of an information element and a particular one of a plurality of information element segments for writing to the selected one of the plurality of memory channels.

42. (new) The system of claim 41 further comprising

a data buffer unit to temporarily store at least one of the information element and the particular one of the plurality of information element segments within a particular one of a plurality of write payload channel queues that corresponds to the selected one of the plurality of memory channels and writes the temporarily stored data to the selected one of the plurality of memory channels within the memory unit; and

a forward processing unit that fetches at least one buffer of the plurality of buffers within the memory unit.

43. (new) The system of claim 42 wherein the forward processing unit includes

a plurality of read payload channel queues, each of the plurality of read payload channel queues corresponds to a separate one of the plurality of memory channels, each of the plurality of read payload channel queues stores at least one of an information element and a particular one of the information element segments that is retrieved from the memory unit; and

a plurality of read request queues, each of the plurality of read request queues corresponds to a separate one of the plurality of read payload channel queues, a particular one of the plurality of read request queues stores a request to retrieve from a corresponding one of the plurality of memory channels of the memory unit at least one of the information element and the particular one of the information element segments and store it in a corresponding one of the plurality of read payload channel queues.

44. (new) The system of claim 43 further comprising a payload channel sequence table to specify an at least one memory channel of the plurality of memory channels at which at least one of the information element and a portion of the information element is stored.

45. (new) The system of claim 43 wherein the forward processing unit fetches at least one of the information element and the portion of the information element by

accessing the payload channel sequence table to determine at least one memory channel within which at least one of the information element and the portion of the information element is stored, and

for each of the at least one memory channel, sending a read request to a particular one of the plurality of read request queues that corresponds to that memory channel.

46. (new) The system of claim 44 wherein the payload channel sequence table is partitioned to form a plurality of columns, each of the plurality of columns corresponds to a separate one of the plurality of memory channels, the plurality of columns are further partitioned to form a plurality of rows, each of the plurality of rows includes a plurality of buffer information units and each of the plurality of buffer information units includes

a buffer location field that specifies a particular one of the plurality of memory channels at which a particular one of the plurality of buffers at a particular one of the plurality of memory lines stores at least one of the information element and a particular one of the plurality of information element segments;

a value field that indicates whether the particular one of the plurality of buffers corresponding to the particular one of the plurality of memory channels at the particular one of the plurality of memory lines stores any data within that buffer; and

an end-of-packet field that indicates whether the particular one of the plurality of buffers corresponding to the particular one of the plurality of memory channels at the particular one of the plurality of memory lines stores data that signals an end-of-packet.

47. (new) The system of claim 46 wherein the forward processing unit fetches at least one of the information element and the portion of the information element by

determining at least one memory channel that stores the at least one of the information element and the portion of the information element by traversing each of the plurality of buffer information units within a particular one of the plurality of rows of the payload channel sequence table and retrieving the particular one of the plurality of memory channels specified within the buffer location field until the end-of-packet field of that buffer information unit signals the end-of-packet; and

for each of the at least one memory channel, sending a read request to a particular one of the plurality of read request queues that corresponds to that memory channel.

48. (new) A program storage device readable by a computer system, storing a plurality of instructions to optimally access a memory unit where the memory unit is logically partitioned to form a plurality of memory channels, the plurality of memory channels are further logically partitioned to form a plurality of memory lines, each of the plurality of memory lines includes a plurality of buffers and each of the plurality of buffers corresponds to a separate one of the plurality of memory channels, comprising:

instructions for determining at least one load value of each of the plurality of memory channels; and

instructions for selecting a particular one of the plurality of memory channels based on the determined at least one load value.

49. (new) The device of claim 48 wherein the instructions for determining the at least one load value of each of the plurality of memory channels includes instructions for determining, for each of the plurality of memory channels, at least one of the number of pending read requests, the number of pending write requests, and the number of active buffers which is the number of a

particular one of the plurality of buffers that is unavailable and corresponds to the particular one of the plurality of memory channels in each of the plurality of memory lines.

50. (new) The device of claim 48 wherein the instructions for selecting the particular one of the plurality of memory channels includes instructions for selecting the particular one of the plurality of memory channels that has at least one of a lowest number of pending read requests, a lowest number of pending write requests, a lowest number of active buffers, and a corresponding channel identification number that is next in a round robin scheme.

REMARKS

A provisional patent application entitled Memory Management System and Algorithm for Network Processor Architecture having the application number 60/242,831 was filed on October 23, 2000. A petition, accompanying this preliminary amendment, requests conversion of that provisional application to a nonprovisional application pursuant to 37 C.F.R. § 1.53(c)(3). This preliminary amendment adds claims 1-50 to the resulting nonprovisional application. Support for new claims 1-50 can be found throughout the specification. No new matter has been added.

Respectfully submitted,

Dated: October 22, 2001 By: Thomas George

Thomas George

Registration No.: 45,740

Morrison & Foerster LLP
755 Page Mill Road
Palo Alto, California 94304-1018
Telephone: (650) 813-5722
Facsimile: (650) 494-0792